

SPECIFICATION

Electronic Version 1.2.8

Stylesheet Version 1.0

Method and System for Optimizing Data Searches in Tree Structures

Field of the Invention

[0001] The present invention relates to control structures for tree searches in embedded processing systems.

Background of the Invention

[0002] Processing system designers continually seek new ways to improve device performance. While processing speeds continue to increase, the latency imposed by memory access times imposes operating delays. In systems-on-a-chip/embedded systems, efforts to avoid such latency issues have included utilizing local memory in the form of SRAM (static random access memory) on-chip. However, cost and size limitations reduce the effectiveness of the use of SRAM on-chip for some processing environments.

[0003] For example, currently in network environments, network switches are being used to perform more complex operations than simple packet forwarding. Network processors are being developed to provide for more complex processing in network routers, while maintaining flexibility to accommodate changes and enhancements to the functionality provided by the routers, as techniques and protocols evolve. As with most any form of processors, these network processors also face challenges in terms of memory utilization, particularly due to the need to handle a vast array of network traffic.

[0004]

In embedded processing systems, such as network processors, off-chip/external DRAM (dynamic random access memory) is an option that is often chosen due to its lower cost, as compared with SRAM. Thus, while potentially most cost effective, the

use of external DRAM introduces a performance penalty in the form of longer access latency (additional delay cycles for the first request for data) relative to other types of RAM. Further, the problem of longer access latency is felt more sharply with shared DRAM, which needs to support concurrent operations required by the system, such as reading in new data from a DMU (data management unit) at the same time that a search for data in the memory is being performed.

[0005] In order to facilitate quicker storage and retrieval of data from the DRAM, a tree structure often is employed for the data being stored. For example, a typical tree structure may be from 12 levels to more than 23 levels deep. Such a large number of levels requires multiple requests to memory to obtain all of the necessary data, i.e., to access and utilize the desired leaf of the tree. In addition, with each successive level of the tree, there is more data (unsearched) than the previous level. These factors create further issues regarding how quickly traversal of a tree structure can occur.

[0006] Accordingly, what is needed is a system and method for optimization of a tree structure for data stored in external DRAM of an embedded processing system. The present invention addresses such a need.

Brief Summary of the Invention

[0007] Aspects for optimizing data searches in tree structures are described. The aspects include organizing multiple search levels of data into sub-trees contained in fixed size blocks of shared external memory of an embedded processing system, and requiring each reference to the data to proceed from one-half of a sub-tree during a descent of the search tree based on a search pattern.

[0008] With the organization of PSCBs in a tree structure in accordance with the present invention, optimization of memory latency while descending levels of tree is achieved, since a larger piece of data is referenced and used more than once during descent of the tree, with local subsections of the tree in one piece of memory. In this manner, faster search operations on large tree structures can be realized, which aids in alleviating latency issues that utilization of external, shared memory impose in embedded processing systems. These and other advantages of the present invention will be more fully understood in conjunction with the following detailed description

and accompanying drawings.

Brief Description of the Several Views of the Drawings

- [0009] Figure 1 illustrates an overall block diagram of an embedded processing system.
- [0010] Figure 2 shows a table reflecting optimization of FM/SM group size based on an example of a 12 level tree.
- [0011] Figures 3a and 3b present graphs for the results of evaluating the performance and resource usage for a wide range of possible tree depths (1 through 30) for FM and SM searches.
- [0012] Figure 3c presents a graph of the overall per level average or slope of the curves performance and resource usage Graphs 1 and 2 of Figures 3a and 3b.
- [0013] Figure 4 shows a table reflecting optimization of LPM group size based on a 12 level tree.
- [0014] Figure 5 illustrates a search tree structure of PSCBs in accordance with the present invention.
- [0015] Figures 6a, 6b, and 6c illustrate organization of PSCBs for FM, LPM and SMT algorithms in accordance with the present invention.

Detailed Description of the Invention

- [0016] The present invention relates to control structures for tree searches in embedded processing systems. The following description is presented to enable one of ordinary skill in the art to make and use the invention and is provided in the context of a patent application and its requirements. Various modifications to the preferred embodiment and the generic principles and features described herein will be readily apparent to those skilled in the art. Thus, the present invention is not intended to be limited to the embodiment shown but is to be accorded the widest scope consistent with the principles and features described herein.

[0017]

The present invention present aspects of providing optimal performance in a processing system utilizing shared RAM memories for both data and control storage.

An overall block diagram of an embedded processing system applicable for utilization of the present invention is illustrated in Figure 1. As shown, the system 10 includes a central processing unit (CPU) core 12, the CPU core including a CPU 14, a memory management unit (MMU) 16, an instruction cache (I-cache) 18, and data cache (D-cache) 20, as is well appreciated by those skilled in the art. A processor local bus 22 couples the CPU core 12 to on-chip SRAM 24. Further coupled to the bus 22 is SDRAM (synchronous DRAM) controller 26, which is coupled to off-chip/external SDRAM 28. A PCI (peripheral component interconnect) bridge 30 is also coupled to bus 22, the PCI bridge 30 further coupled to a host bus 32 that is coupled to host memory 34. As shown, a tree search engine 36 is also included and coupled to bus 22. The tree search engine 36 is a hardware assist that performs pattern analysis through tree searches to find the address of a leaf page for read and write accesses in the SDRAM 28.

[0018] In accordance with the present invention, the searches performed by the tree search engine 36 are improved with the optimization of a tree structure for data stored in external DRAM 28 of an embedded processing system. In general, tree searches, retrievals, inserts, and deletes are performed according to a key. Information is stored in the tree in leaves, which contain the keys as a reference pattern. To locate a leaf, a search algorithm processes input parameters that include the key pattern, and then accesses a direct table (DT) to initiate the walking of the tree structure through pattern search control blocks (PSCBs). The searches occur based on a full match (FM) algorithm, a longest prefix match (LPM) algorithm, or a software management tree (SM) algorithm. The present invention provides a tree structure of PCSBs optimized for all three types of search algorithms, as described hereinbelow.

[0019] An optimization of a tree structure in accordance with the present invention is provided by organizing multiple search levels into sub-trees of PSCBs contained in fixed size blocks of memory and requiring only the left or right side of each sub-tree during each descent of the search tree with the choice of left or right known before the reference of each sub-tree to reduce the size of the required reference.

[0020] Preferably several parameters are considered in determining the organization, including: a latency per reference number, which is determined by the latency of a

memory burst reference plus an adjustment for the expected average bank busy delays; a bus time per reference number, which is determined by the number of memory data bus cycles needed for a memory burst reference plus again the same adjustment for the expected average bank busy delay; a tree search efficiency percentage, which is a metric of the relative per clock search efficiency; a tree search performance number, which is a calculation of the search time portion of the total table lookup performance; a memory bus efficiency percentage, which is a metric of the relative efficiency of the bus usage during the search portion of the table lookup process; and a memory bus resource number, which is a calculation of the memory resources used during the search portion of the table lookup process. The determination of the organization according to the parameters occurs via the following equations for FM and SM:

$$\text{References} = \text{Levels (e.g., 12)} / \text{Levels_per_Reference}$$

$$\text{Tree_Search_Efficiency} = \text{Levels_per_Reference} / \text{Latency_per_Reference}$$

$$\text{Tree_Search_Performance} = \text{References} * \text{Latency_per_Reference}$$

$$\text{FM_SM_Usage} = \text{FM_SM_PSCB_Size (e.g., 4)} / \text{Effective_Bus_Size (e.g., 8)}$$

$$\text{Memory_Bus_Efficiency} = (\text{Levels_per_Reference} / \text{Bus_Time_per_Reference}) * \text{FM_SM_Usage}$$

$$\text{Memory_Bus_Resource} = \text{References} * \text{Bus_Time_per_Reference}$$

[0021] Table 1 in Figure 2 shows optimization of FM/SM group size based on an example of a 12 level tree. The optimum solution was developed in two parts, the first being the organization of multiple search levels into sub-trees contained in fixed size blocks of memory. The second part was the observation that only the left or right side of each sub-tree is required during each decent of the search tree and the choice of left or right is known before the reference of each sub-tree thus reducing the size of the required reference. The 3, 7, and 15 PSCB cases are based on the first part of the solution only, and the 3.5 and 7.5 cases are the extensions of the 7 and 15 PCSB cases respectively based on the second part of the solution.

[0022] The results of evaluating the performance and resource usage for a wide range of possible tree depths (1 through 30) for FM and SM searches can be seen in Graphs 1 and 2 shown in Figure 3a and 3b, respectively. It can be seen in the graphs that the 3.5 PCSBs (plot line 40) and the 7.5 PCSBs (plot line 42) cases are better in both

performance and resource usage than all the other FM and SM solutions and show significant improvements over the 1 PSCB (plot line 44) control case. Plot line 46 shows the 3 PSCBs case, plot line 48 shows the 7 PSCBs case, and plot line 50 shows the 15 PSCBs case. As appears from Graph 1, the performance of the 7.5 case is better at most depths than the 3.5 case but that in Graph 2, the resource usage of the 3.5 case is better at most depths than the 7.5 case.

[0023] The overall per level average or slope of the curves performance and resource usage Graphs 1 and 2 can be seen in Graph 3 in Figure 3c. As shown in Graph 3, the resource minimum is at the 3.5 PSCB's point (node 52) and the performance maximum (clock minimum) is at the 7.5 PSCBs point (node 54). The performance difference between the two points is 0.5 clocks per level and the resource difference is 0.4 cycles per level. It has been found that a five tenths of a clock improvement in performance is worth the four tenths of a cycle increase in resource usage making the 7.5 PSCB's case the optimum solution for FM and SM tree searches in the example embodiment.

[0024] Table 2 of Figure 4 shows optimization of LPM group size based on 12 level tree. The 1.5 and 3.5 cases are based on both parts of the solution and require similar block sizes as the 3.5 and 7.5 cases of Table 1 of Figure 2, respectively. The memory bus efficiency calculation has changed to reflect the difference in the LPM PSCB size:

$$\text{LPM_Usage} = \text{LPM_PSCB_Size (e.g., 7)} / \text{Effective_Bus_Size (e.g., 8)}$$

$$\text{Memory_Bus_Efficiency} = (\text{Levels_per_Reference} / \text{Bus_Time_per_Reference}) \\ * \text{LPM_Usage}$$

[0025] As can be seen in Table 2, the 3.5 solution has the best performance while having the same resource usage as the 1.5 case making the 3.5 solution the optimum solution for LPM tree searches with the added benefit of using the same block size as the 7.5 PSCB's case from the FM and SM tree search solution.

[0026] Figure 5 illustrates a search tree structure of PSCBs in accordance with the present invention. By way of example, a search of the tree in Figure 5 begins with the memory access request of the left or right half of the Root or level 0 Branch Table (BT) based on the Next Bit Test (NBT) result from the Lookup Definition (LUDef) or Direct Table (DT, not shown) entry for this search tree. The access of the first branch table half contains the optimum number of levels of PSCBs of the tree for the search type. If

after descending through the first table an external (lower) branch table address is arrived at instead of a leaf address, then an additional memory access request would be made for only the left or right half of this lower branch table. This process continues until a leaf address is arrived at during the decent through the lower branch table halves. When the search arrives at a leaf address, the process terminates with a memory access request for the leaf data to determine if a match was found. The leaf structure for the leaves shown in Figure 5 is described more particularly in co-pending U.S. Patent Application, filed ____, serial no. ____ (docket no. RPS920020019US1/2492P), assigned to the assignee of the present invention, and incorporated herein by reference in its entirety.

[0027] Representations of a basic organization of PSCBs for each type of search algorithm, FM, LPM, and SM, are illustrated in Figures 6a, 6b, and 6c. In each of these figures, local pointer (LP) values provide reference to a table within the retrieved data, and traversal based on a LP is illustrated by the dashed arrows. Branch table (BT) values provide reference to a table outside of the retrieved data, and traversal based on a BT value is illustrated by the solid arrows. EXP value provide expiration data for found LEAF data.

[0028] With the organization of PSCBs in a tree structure in accordance with the present invention, optimization of memory latency while descending levels of tree is achieved, since a larger piece of data is referenced and used more than once during descent of the tree, with local subsections of the tree in one piece of memory. In this manner, faster search operations on large tree structures can be realized, which aids in alleviating latency issues that utilization of external, shared memory impose in embedded processing systems.

[0029] Although the present invention has been described in accordance with the embodiments shown, one of ordinary skill in the art will readily recognize that there could be variations to the embodiments and those variations would be within the spirit and scope of the present invention. Accordingly, many modifications may be made by one of ordinary skill in the art without departing from the spirit and scope of the appended claims.